



UNITED STATES PATENT AND TRADEMARK OFFICE

[Signature]
UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/608,037	06/30/2003	Sanjay Ghemawat	0026-0031	8255
44989	7590	01/25/2006	EXAMINER	
HARRITY SNYDER, LLP 11350 Random Hills Road SUITE 600 FAIRFAX, VA 22030			LOVEL, KIMBERLY M	
			ART UNIT	PAPER NUMBER
			2167	

DATE MAILED: 01/25/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)	
	10/608,037	GHEMAWAT ET AL.	
	Examiner	Art Unit	
	Kimberly Lovel	2167	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 30 June 2003.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-18 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-18 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☒ The drawing(s) filed on 30 June 2003 is/are: a) ☒ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|--|---|
| 1) <input checked="" type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413) |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | Paper No(s)/Mail Date. _____ |
| 3) <input checked="" type="checkbox"/> Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08) | 5) <input type="checkbox"/> Notice of Informal Patent Application (PTO-152) |
| Paper No(s)/Mail Date <u>12/03/03 2/03/2004</u> . | 6) <input type="checkbox"/> Other: _____ |

DETAILED ACTION

Claim Status

1. Claims 1-18 are pending.
2. Claims 1-18 are rejected.

Information Disclosure Statement

3. The information disclosure statement filed 30 June 2003 fails to comply with 37 CFR 1.98(a)(2), which requires a legible copy of each cited foreign patent document; each non-patent literature publication or that portion which caused it to be listed; and all other information or that portion which caused it to be listed. Copies of Applications 10/608,139, 10/608,039, 10/608,135 and 10/608,136 were not received. However, the information referred to therein has been considered since the applications are US applications and can be viewed in the USPTO's internal system.

Claim Rejections - 35 USC § 102

4. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this

Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

Art Unit: 2167

5. Claims 1-7 and 12-18 are rejected under 35 U.S.C. 102(b) as being anticipated by US Patent No. 5,991,804 issued to Bolosky et al (hereafter Bolosky et al).

Claim 1:

Referring to claim 1, Bolosky et al disclose a file system (see abstract and Fig 1, item 20), comprising:

a plurality of servers configured to store data (see column 1, line 67 – column 2, line 1 – the data servers are considered to represent the plurality of *servers configured to store data*; column 3, lines 47-50; Fig 1, items 24 and 28 – item 24 represents the data server which contains item 28 which represents a data disk that stores data; column 7, lines 20-35; and Fig 4); and

a master connected to the servers (see column 1, lines 66-67 – the controller is considered to represent the *master*; column 3, lines 43-45; Fig 1, items 22, 24 and 26; column 7, lines 4-19; and Fig 3) and configured to:

communicate with the servers upon startup of the master to authoritatively identify the data stored by the servers (see column 2, lines 55-57 and column 16, lines 34-48 – during whole system failure, it can be concluded that the controller fails; when the system restarts, the controller uses the log to determine the data stored by the servers in order to continue the transfer process), and

record location information that identifies ones of the servers that store the data (see column 5, lines 1-3 – the map tracks the location of the data block which identifies which server the data is stored on).

Art Unit: 2167

Claim 2:

Referring to claim 2, Bolosky et al disclose the system of claim 1, wherein the data corresponds to files stored as chunks by the servers (see column 4, lines 37-44 – the disk stores data in blocks which is considered to represent a chunk since a block is defined as a unit of data or an amount of physical space allocated on a disk to hold that unit of data).

Claim 3:

Referring to claim 3, Bolosky et al disclose the system of claim 1, wherein the master is further configured to control placement of new data at the servers (see column 8, lines 13-16 – when performing striping, the controller acts as if it were writing the data files for the first time which is considered to represent writing new data).

Claim 4:

Referring to claim 4, Bolosky et al disclose the system of claim 3, wherein when controlling the placement of new data, the master is configured to:

identify one or more of the servers to store the new data based on at least one of utilization of the servers (column 7, line 66 – column 8, line 2), prior chunk distribution involving the servers (see column 8, lines 31-42), network topology, and failure correlation properties associated with the servers, and

place the data at the identified one or more servers (see column 8, lines 16-22).

Claim 5:

Art Unit: 2167

Referring to claim 5, Bolosky et al disclose the system of claim 1, wherein the master is further configured to control redistribution of the data stored by the servers (see column 6, lines 30-38).

Claim 6:

Referring to claim 6, Bolosky et al disclose the system of claim 5, wherein when controlling redistribution of the data, the master is configured to:

select data to redistribute based on a current distribution of the data (see column 8, lines 11-22),

identify on one or more of the servers to which to move the selected data (see column 8, lines 31-53), and

move the selected data to the identified one or more servers (see column 9, line 64 – column 10, line 14).

Claim 7:

Referring to claim 7, Bolosky et al disclose the system of claim 1, wherein the master is further configured to monitor the state of the servers (see column 7, line 10; Fig 1, item 46; and column 16, lines 7-33 – the data server coordinating program allows for the controller to monitor the state of the servers).

Claim 12:

Referring to claim 12, Bolosky et al disclose the system of claim 1, wherein the location information is not stored persistently by the master (see column 4, line 64 – column 5, line 3 – the system provides two alternatives for the location information; the master does not persistently store the location

Art Unit: 2167

information in the first alternative mentioned in column 4, line 66 – column 5, line 1).

Claim 13:

Referring to claim 13, Bolosky et al disclose in a file system that includes a master connected to a plurality of servers (see abstract and Fig 1), the master comprising:

means for performing a startup operation (see column 2, lines 55-57 – the controller represents the means; column 7, lines 4-19; and column 16, lines 34-48);

means for communicating with the servers during or after the startup operation to authoritatively identify the data stored by the servers (see column 3, lines 43-45 and Fig 1, item 26 – the low bandwidth control network represents the means); and

means for storing location information that identifies ones of the servers that stores the data (see column 5, lines 1-3 – the map represents the means).

Claim 14:

Referring to claim 14, Bolosky et al disclose a method for maintaining data in a file system that includes a master connected to a plurality of servers, the method, performed by the master (see Fig 1), comprising:

communicating with the servers upon startup of the master to authoritatively determine data stored by the servers (see column 2, lines 55-57 and column 16, lines 34-48 – during whole system failure, it can be concluded that the controller fails; when the system restarts, the controller uses the log to

Art Unit: 2167

determine the data stored by the servers in order to continue the transfer process); and

generating location information based on the data determined to be stored by the servers (see column 5, lines 1-3 – the map tracks the location of the data block which identifies which server the data is stored on).

Claim 15:

Referring to claim 15, Bolosky et al disclose a file system (see abstract and Fig 1, item 20), comprising:

a plurality of servers configured to store files as chunks (see column 1, line 66 – column 2, line 3 and column 4, lines 37-44 – the disk stores data in blocks which is considered to represent a chunk since a block is defined as a unit of data or an amount of physical space allocated on a disk to hold that unit of data); and

a master connected to the servers (see column 1, lines 66-67 – the controller is considered to represent the *master*; column 3, lines 43-45; Fig 1, items 22, 24 and 26; column 7, lines 4-19; and Fig 3) and configured to:

authoritatively determine location information by communicating with the servers, the location information being based on which of the servers stores ones of the chunks (see column 3, lines 43-45 – the low bandwidth control network provides the communication and column 5, lines 1-3 – the map tracks the location), and

Art Unit: 2167

update the location information by periodically communicating with the servers to obtain changes to the location information (see column 9, line 64 – column 10, line 14).

Claim 16:

Referring to claim 16, Bolosky et al disclose a file system (see abstract and Fig 1, item 20), comprising:

a plurality servers configured to store files as chunks (see column 1, line 66 – column 2, line 3 and column 4, lines 37-44 – the disk stores data in blocks which is considered to represent a chunk since a block is defined as a unit of data or an amount of physical space allocated on a disk to hold that unit of data); and

a master connected to the servers (see column 1, lines 66-67 – the controller is considered to represent the *master*; column 3, lines 43-45; Fig 1, items 22, 24 and 26; column 7, lines 4-19; and Fig 3) and configured to:

communicate with the servers to authoritatively determine location information of the data, the location information being based on which of the servers store the data (see column 3, lines 43-45 – the low bandwidth control network provides the communication and column 5, lines 1-3 – the map tracks the location),

periodically communicate with the servers to obtain changes to the location information (see column 9, line 64 – column 10, line 14), and

update the location information based on the changes to the location information (see column 9, line 64 – column 10, line 14).

Art Unit: 2167

Claim 17:

Referring to claim 17, Bolosky et al disclose A file system (see abstract and Fig 1, item 20), comprising:

a plurality of servers configured to store data (see column 1, line 67 – column 2, line 1 – the data servers are considered to represent the plurality of *servers configured to store data*; column 3, lines 47-50; Fig 1, items 24 and 28 – item 24 represents the data server which contains item 28 which represents a data disk that stores data; column 7, lines 20-35; and Fig 4); and

a master connected to the servers (see column 1, lines 66-67 – the controller is considered to represent the *master*; column 3, lines 43-45; Fig 1, items 22, 24 and 26; column 7, lines 4-19; and Fig 3) and configured to:

communicate with the servers to authoritatively determine location information of the data, the location information being based on which of the servers store the data (see column 3, lines 43-45 – the low bandwidth control network provides the communication and column 5, lines 1-3 – the map tracks the location),

instruct one of the servers to perform an action concerning the data, the action causing a change in the location information (see column 9, line 44 – column 10, line 14 – occurs during restriping), and

update the location information based on the change to the location information upon completion of the action (see column 9, line 64 – column 10, line 14) .

Claim 18:

Art Unit: 2167

Referring to claim 18, Bolosky et al disclose a file system (see abstract and Fig 1, item 20), comprising:

a plurality of servers configured to store data (see column 1, line 67 – column 2, line 1 – the data servers are considered to represent the plurality of *servers configured to store data*; column 3, lines 47-50; Fig 1, items 24 and 28 – item 24 represents the data server which contains item 28 which represents a data disk that stores data; column 7, lines 20-35; and Fig 4); and

a master connected to the servers (see column 1, lines 66-67 – the controller is considered to represent the *master*; column 3, lines 43-45; Fig 1, items 22, 24 and 26; column 7, lines 4-19; and Fig 3) and configured to:

communicate with the servers to authoritatively determine information regarding the data (see column 3, lines 43-45 – the low bandwidth control network provides the communication and column 5, lines 1-3 – the map tracks the location),

instruct one of the servers to perform an action concerning the data, the action causing a state change associated with the information (see column 9, line 44 – column 10, line 14 – the change in location of the data represents the state change), and

update state information based on the state change upon completion of the action (see column 9, line 64 – column 10, line 14).

Claim Rejections - 35 USC § 103

6. This application currently names joint inventors. In considering patentability of the claims under 35 U.S.C. 103(a), the examiner presumes that the subject matter of the various claims was commonly owned at the time any inventions covered therein were made absent any evidence to the contrary. Applicant is advised of the obligation under 37 CFR 1.56 to point out the inventor and invention dates of each claim that was not commonly owned at the time a later invention was made in order for the examiner to consider the applicability of 35 U.S.C. 103(c) and potential 35 U.S.C. 102(e), (f) or (g) prior art under 35 U.S.C. 103(a).

7. Claims 8-9 are rejected under 35 U.S.C. 103(a) as being unpatentable over US Patent No. 5,991,804 issued to Bolosky et al as applied to claim 7 above, and further in view of US Patent No 6,725,392 issued to Frey et al (hereafter Frey et al).

Bolosky et al disclose a system for providing a master that is further configured to monitor the state of the servers as taught above. However, Bolosky does not explicitly teach the further limitation of the master being configured to exchange heartbeat signals with the servers to determine the state of the servers. Frey et al disclose a system similar to that of Bolosky, including the master being configured to exchange heartbeat signals with the servers to determine the state of the servers.

Claim 8:

Art Unit: 2167

In particular, Frey et al teach a system similar to that of claim 7, wherein the master is configured to exchange heartbeat signals with the servers to determine the state of the servers (see column 14, lines 15-19 – the state is whether or not the server is operational).

It would have been obvious to one of ordinary skill in the art at the time the invention was made to add Frey et al's use of heartbeat signals to Bolosky et al's system for monitoring the state of the servers. One would have been motivated to do so since both Frey et al and Bolosky et al deal with fault recovery systems for distributed file systems having a controller (Frey et al: see abstract; Bolosky et al: see abstract).

Claim 9:

Referring to claim 9, the combination of Frey et al and Bolosky et al discloses the system of claim 8, wherein the heartbeat signals include space utilization information (see column 14, lines 15-29 – memory access request is considered to represent space utilization information).

8. Claims 10-11 are rejected under 35 U.S.C. 103(a) as being unpatentable over US Patent No. 5,991,804 issued to Bolosky et al as applied to claim 7 above, and further in view of US PGPub 2003/0023898 issued to Jacobs et al (hereafter Jacobs et al).

Bolosky et al disclose a system for providing a master that is further configured to monitor the state of the servers as taught above. However, Bolosky does not explicitly teach the further limitation wherein the state of the

Art Unit: 2167

servers includes information regarding the data stored by the servers. Jacobs et al disclose a system similar to that of Bolosky, including the state of the servers including information regarding the data stored by the servers.

Claim 10:

In particular, Jacobs et al teach a system similar to that of claim 7, wherein the state of the servers includes information regarding the data stored by the servers (see paragraph [0009], lines 6-10).

It would have been obvious to one of ordinary skill in the art at the time the invention was made to include Jacobs et al's information regarding the data stored by the servers with Bolosky et al's system for monitoring the state of the servers. One would have been motivated to do so since both Jacobs et al and Bolosky et al deal with sending versions of data from a master server to a second server (Frey et al: see abstract; Bolosky et al: see abstract).

Claim 11:

Referring to claim 11, the combination of Jacobs et al and Bolosky et al discloses the system of claim 10, wherein the information includes version numbers of the data (see paragraph [0009], lines 6-10).

Conclusion

9. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure.

Art Unit: 2167

- US Patent No 6,889,249 issued to Miloushev et al, which discusses a file system using a file switch computer that performs similar functions to that of the applicant's master.
- US Patent No 6,134,596 issued to Bolosky et al, which includes further discussion of the media file server system utilized by Bolosky et al in US Patent No 5,991,804.
- US Patent No 6,173,293 issued to Thekkah et al. The patent discusses the Frangipani File System, which is scalable and distributed.

10. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Kimberly Lovel whose telephone number is (571) 272-2750. The examiner can normally be reached on M-F from 8:00 to 3:30.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Jean Homere, can be reached on (571) 272-3780. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

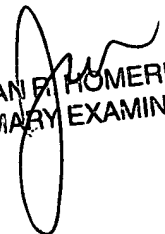
Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR

Art Unit: 2167

system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Kimberly Lovel
Examiner
Art Unit 2167

kml
19 January 2006


JEAN P. HOMERE
PRIMARY EXAMINER

BC